




US 20130246446A1

(19) **United States**(12) **Patent Application Publication**
Taranov(10) **Pub. No.: US 2013/0246446 A1**(43) **Pub. Date: Sep. 19, 2013**(54) **OPTIMIZING DATA PROCESSING USING
DYNAMIC SCHEMAS**(71) Applicant: **MICROSOFT CORPORATION**,
Redmond, WA (US)(72) Inventor: **Igor Taranov**, Bellevue, WA (US)(73) Assignee: **Microsoft Corporation**, Redmond, WA
(US)(21) Appl. No.: **13/887,557**(22) Filed: **May 6, 2013****Related U.S. Application Data**(63) Continuation of application No. 13/180,131, filed on
Jul. 11, 2011, now Pat. No. 8,458,203.**Publication Classification**(51) **Int. Cl.**
G06F 17/30 (2006.01)(52) **U.S. Cl.**CPC **G06F 17/30569** (2013.01)USPC **707/756**(57) **ABSTRACT**

A computer system accesses rows of feed data and converts the received feed data into portions of binary blob data. The computer system also sends the binary blob data to a database server which is configured to access metadata associated with a feed including a dynamic server statement to determine how to convert the binary blob data to a server table with a blob column configured to store the rows of feed data. The database server accesses feed data belonging to a particular feed and executes a dynamic server statement to create a relational dataset in an in-memory table of the server. A second dynamic statement applies data processing conditions indicated in the metadata. When feed data rows match conditions, the computer system places feed data row information into an alert table that includes references to the blob table with blob data, thereby triggering an alert.






Figure 1

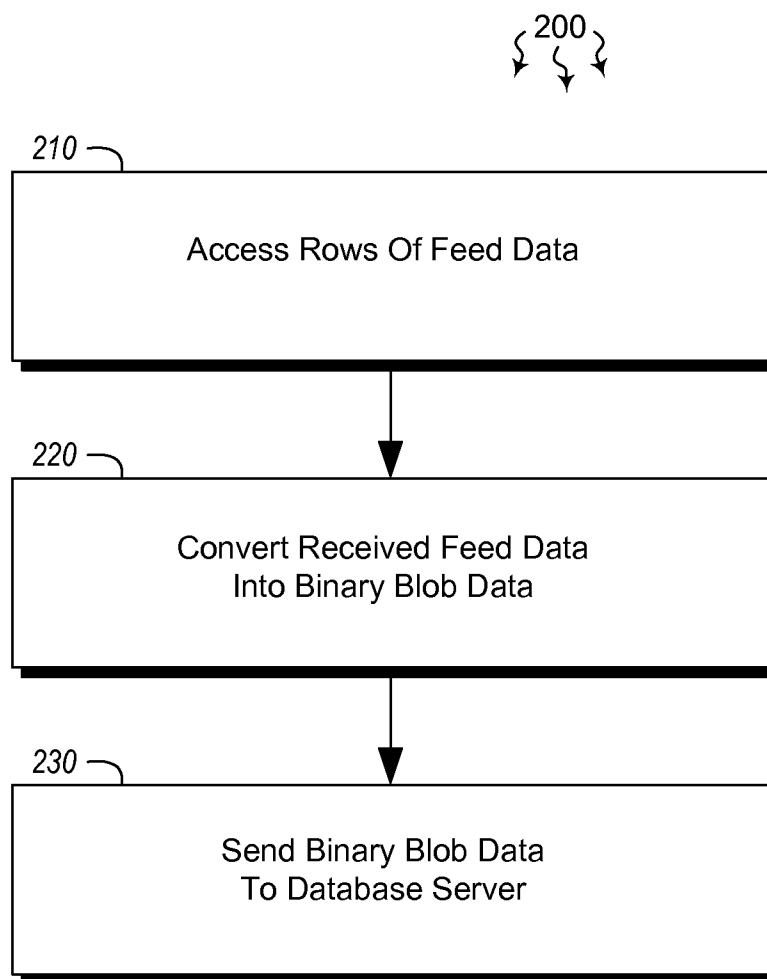
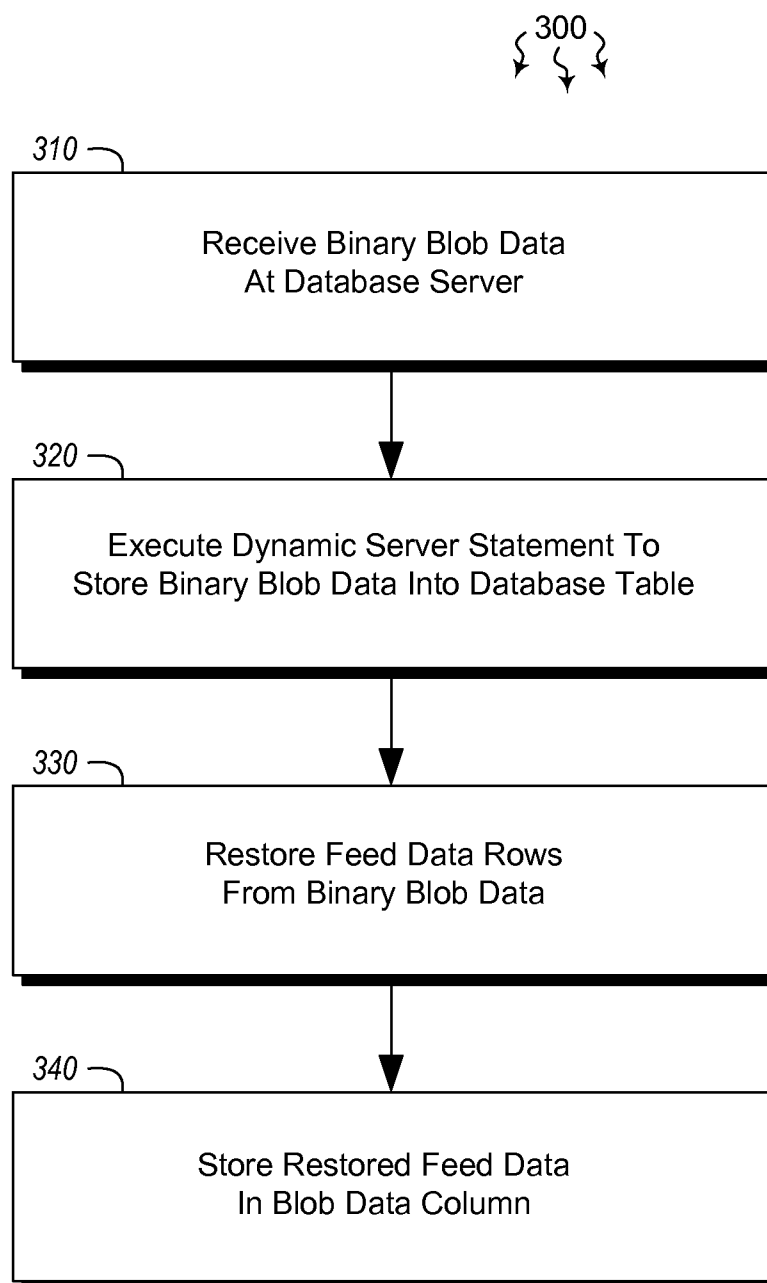
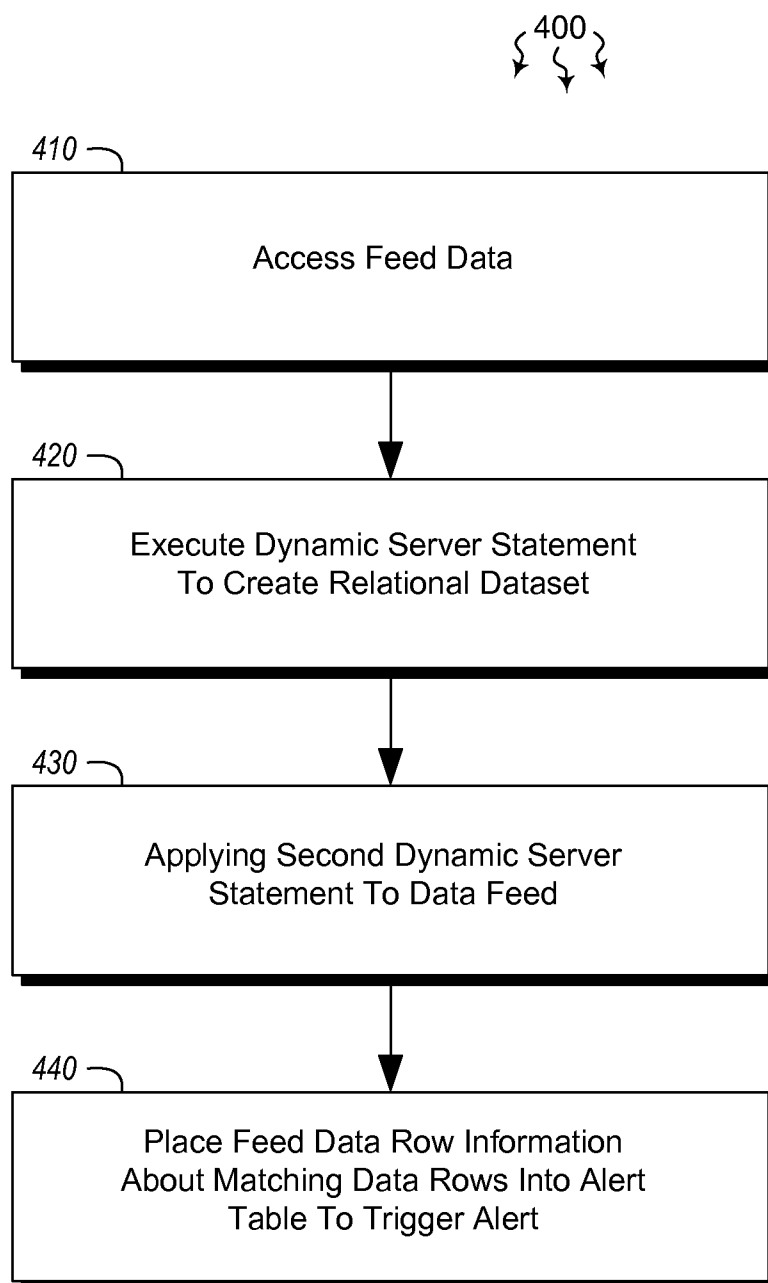


Figure 2

**Figure 3**

**Figure 4**




Figure 5

OPTIMIZING DATA PROCESSING USING DYNAMIC SCHEMAS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation of U.S. patent application Ser. No. 13/180,131, filed Jul. 11, 2011, and entitled "Optimizing Data Processing Using Dynamic Schemas." The contents of the foregoing application are incorporated by reference herein in their entirety.

BACKGROUND

[0002] Computers have become highly integrated in the workforce, in the home, in mobile devices, and many other places. Computers can process massive amounts of information quickly and efficiently. Software applications designed to run on computer systems allow users to perform a wide variety of functions including business applications, school-work, entertainment and more. Software applications are often designed to perform specific tasks, such as word processor applications for drafting documents, or email programs for sending, receiving and organizing email.

[0003] In many cases, software applications are designed to interact with other software applications or other computer systems. For example, applications may communicate with databases to store data. These databases may receive the data and store and process it according to predefined, static data schemas. These data schemas are specific to each type of data, and only a certain number of static, predefined schemas can be used in a given database implementation.

BRIEF SUMMARY

[0004] Embodiments described herein are directed to efficiently uploading data to a database using dynamic schemas and to efficiently processing data using dynamic schemas. In one embodiment, a computer system accesses rows of feed data. The feed data includes data fields of fixed and varying lengths. The computer system converts the received feed data into portions of binary blob data. The binary blob data is allocated in fixed-size data portions representing feed rows. Varying length feed fields are represented by both a fixed-size field, with a length of the feed field, and a reference into the blob data that includes the actual varying size feed data. The computer system also sends the binary blob data to a database server. The database server is configured to access metadata associated with a feed. Metadata contains dynamic server statement to determine how to convert the binary blob data to a server table with a BLOB column configured to store the rows of feed data.

[0005] In another embodiment, a database server receives from an application server various portions of binary blob data. The database server executes a dynamic server statement to store the binary blob data into a database table. The dynamic server statement is stored within metadata associated with the feed. The database server restores feed data rows from the binary blob data by substituting varying size references with the actual feed data, while keeping the feed data in a binary format. The database server also stores the restored feed data in the database table in a BLOB data column. Each feed data row results in one binary blob value in one corresponding row in the database table.

[0006] In yet another embodiment, a database server accesses feed data belonging to a particular feed and executes

a dynamic server statement to create a relational dataset with data type fields from the feed data in an in-memory table of the server. The dynamic server statement is stored within metadata associated with the feed. Then database server applies a second dynamic server statement to the data feed which applies data processing conditions indicated in the metadata. Also, upon determining that one or more feed data rows match the data processing conditions the database server places feed data row information about the matching data rows into an alert table that includes references to the table with the BLOB column containing the feed data, thereby triggering an alert.

[0007] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

[0008] Additional features and advantages will be set forth in the description which follows, and in part will be obvious from the description, or may be learned by the practice of the teachings herein. Features and advantages of the invention may be realized and obtained by means of the instruments and combinations particularly pointed out in the appended claims. Features of the present invention will become more fully apparent from the following description and appended claims, or may be learned by the practice of the invention as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] To further clarify the above and other advantages and features of embodiments of the present invention, a more particular description of embodiments of the present invention will be rendered by reference to the appended drawings. It is appreciated that these drawings depict only typical embodiments of the invention and are therefore not to be considered limiting of its scope. The invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

[0010] FIG. 1 illustrates a system architecture in which embodiments of the present invention may operate including efficiently uploading data to a database using dynamic schemas.

[0011] FIG. 2 illustrates an application server side flowchart of an example method for efficiently uploading data to a database using dynamic schemas.

[0012] FIG. 3 illustrates a database server side flowchart of an example method for efficiently uploading data to a database using dynamic schemas.

[0013] FIG. 4 illustrates a database server side flowchart of an example method for efficiently processing data using dynamic schemas.

[0014] FIG. 5 illustrates a database server side example application of dynamic server statements for efficiently processing data using dynamic schemas.

DETAILED DESCRIPTION

[0015] Embodiments described herein are directed to efficiently uploading data to a database using dynamic schemas and to efficiently processing data using dynamic schemas. In one embodiment, a computer system accesses rows of feed data. The feed data includes data fields of fixed and varying lengths. The computer system converts the received feed data

into portions of binary blob data. The binary blob data is allocated in fixed-size data portions representing feed rows. Varying length fields are represented by both a fixed-size field, with a length of the feed field, and a reference into the blob data that includes the actual varying size feed data. The computer system also sends the binary blob data to a database server. The database server is configured to access metadata associated with a feed. Metadata contains dynamic server statement to determine how to convert the binary blob data to a server table with a BLOB column configured to store the rows of feed data.

[0016] In another embodiment, a database server receives from an application server various portions of binary blob data. The database server executes a dynamic server statement to store the binary blob data into a database table. The dynamic server statement is stored within metadata associated with the blob data's original feed. The database server restores feed data rows from the binary blob data by substituting varying size references with the actual feed data, while keeping the feed data in a binary format. The database server also stores the restored feed data in the database table in a BLOB data column. Each feed data row results in one binary blob value in one corresponding row in the database table.

[0017] In yet another embodiment, a database server accesses feed data belonging to a particular feed and executes a dynamic server statement to create a relational dataset with data type fields from the feed data in an in-memory table of the server. The dynamic server statement is stored within metadata associated with the feed. The database server applies a second dynamic server statement to the data feed which applies data processing conditions indicated in the metadata. The dynamic server statement is stored within the metadata associated with the feed. Also, upon determining that one or more feed data rows match the data processing conditions the computer system places feed data row information about the matching data rows into an alert table that includes references to the table with the BLOB column containing the feed data, thereby triggering an alert.

[0018] The following discussion now refers to a number of methods and method acts that may be performed. It should be noted, that although the method acts may be discussed in a certain order or illustrated in a flow chart as occurring in a particular order, no particular ordering is necessarily required unless specifically stated, or required because an act is dependent on another act being completed prior to the act being performed.

[0019] Embodiments of the present invention may comprise or utilize a special purpose or general-purpose computer including computer hardware, such as, for example, one or more processors and system memory, as discussed in greater detail below. Embodiments within the scope of the present invention also include physical and other computer-readable media for carrying or storing computer-executable instructions and/or data structures. Such computer-readable media can be any available media that can be accessed by a general purpose or special purpose computer system. Computer-readable media that store computer-executable instructions in the form of data are computer storage media. Computer-readable media that carry computer-executable instructions are transmission media. Thus, by way of example, and not limitation, embodiments of the invention can comprise at least two distinctly different kinds of computer-readable media: computer storage media and transmission media.

[0020] Computer storage media includes RAM, ROM, EEPROM, CD-ROM, solid state drives (SSDs) that are based on RAM, Flash memory, phase-change memory (PCM), or other types of memory, or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store desired program code means in the form of computer-executable instructions, data or data structures and which can be accessed by a general purpose or special purpose computer.

[0021] A "network" is defined as one or more data links and/or data switches that enable the transport of electronic data between computer systems and/or modules and/or other electronic devices. When information is transferred or provided over a network (either hardwired, wireless, or a combination of hardwired or wireless) to a computer, the computer properly views the connection as a transmission medium. Transmissions media can include a network which can be used to carry data or desired program code means in the form of computer-executable instructions or in the form of data structures and which can be accessed by a general purpose or special purpose computer. Combinations of the above should also be included within the scope of computer-readable media.

[0022] Further, upon reaching various computer system components, program code means in the form of computer-executable instructions or data structures can be transferred automatically from transmission media to computer storage media (or vice versa). For example, computer-executable instructions or data structures received over a network or data link can be buffered in RAM within a network interface module (e.g., a network interface card or "NIC"), and then eventually transferred to computer system RAM and/or to less volatile computer storage media at a computer system. Thus, it should be understood that computer storage media can be included in computer system components that also (or even primarily) utilize transmission media.

[0023] Computer-executable (or computer-interpretable) instructions comprise, for example, instructions which cause a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. The computer executable instructions may be, for example, binaries, intermediate format instructions such as assembly language, or even source code. Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the described features or acts described above. Rather, the described features and acts are disclosed as example forms of implementing the claims.

[0024] Those skilled in the art will appreciate that the invention may be practiced in network computing environments with many types of computer system configurations, including personal computers, desktop computers, laptop computers, message processors, hand-held devices, multiprocessor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, mobile telephones, PDAs, pagers, routers, switches, and the like. The invention may also be practiced in distributed system environments where local and remote computer systems that are linked (either by hardwired data links, wireless data links, or by a combination of hardwired and wireless data links) through a network, each perform tasks (e.g. cloud computing, cloud services and the like). In a

distributed system environment, program modules may be located in both local and remote memory storage devices.

[0025] FIG. 1 illustrates a system architecture **100** in which the principles of the present invention may be employed. System architecture **100** includes an application server system **105** and a database server **125**. These computer systems may be any type of local or distributed computer systems. Various embodiments will be described in relation to these computer systems. For instance, embodiments relating to high efficiency bulk upload of data with dynamic schemas, high efficiency processing of data with dynamic schemas and compact storage of data with dynamic schemas will all be described.

[0026] For instance, a method may be provided that allows uploading of a data of a large volume and with arbitrary schemas by applying the same binary upload approach regardless of the particular schema used. The application server system **105** reads data from an arbitrary feed **101** and packs it in a binary format (blob) **116** as described in FIG. 1. The format allows representing the feed data that might have fields of a varying size by rows with fixed length. Varying size fields are represented by fixed size field data length and references into the second part of the blob, containing the actual varying size data.

[0027] For scalability purposes, large feed data is chunked into pieces according to various specifications (e.g., the number of rows is less than 10000 or chunk size is less than 10 MB, etc.). Each chunk may be uploaded to the database server **125** in a single round trip. The database server reads data and stores it in a table in data store **140**. The database server is aware of the data schema due to the metadata stored for each registered feed. Metadata includes dynamic server statements that the database server can execute dynamically when reading and storing data. The database server side restores the feed data rows by substituting the varying size references with the actual data, but still keeping all data in a binary format. As a result the feed data is stored in the database table in the BLOB column **141**, where each feed row results in one blob value in the corresponding row in the table. It should be noted that field data for fields with varying data size may be different in size for the same field within one feed data reading session.

[0028] Regarding high efficiency processing of data with dynamic schemas, when processing data the database server **125** builds and executes various dynamic server statements, as shown in FIG. 5. Portions of the various dynamic statements are taken from the feed metadata as well. The dynamic statements may be configured to perform the following: reformats incoming feed blob data **116** to represent one row data with one binary value in BLOB column of a database table, reads the feed data **140** belonging to a particular feed reading session from a database table, creates an in-memory database server relational table using various server commands with the definition from the metadata, and applies one or more data processing conditions that are also taken from the metadata. For the rows that satisfy the data conditions, the server side creates references in another database table, pointing to the original database table that contains the binary feed data.

[0029] Regarding compact storage of data with dynamic schemas, the database server **125** may provide compact efficient data storage close in size as if the feed data was stored in a static table. This result may be achieved as the feed data is stored only in the binary format. The metadata is stored per feed registration and so its size is of the same order (i.e. like it would be for a static database table). The system also allows

detection of the data change between subsequent feed reading sessions by hashing each feed row blob value. The database server uses the data change detection mechanism to avoid storing of duplicate data. Each of these concepts will be explained in greater detail below with regard to methods **200**, **300** and **400** of FIGS. 2, 3 and 4, respectively.

[0030] In view of the systems and architectures described above, methodologies that may be implemented in accordance with the disclosed subject matter will be better appreciated with reference to the flow charts of FIGS. 2 and 3. For purposes of simplicity of explanation, the methodologies are shown and described as a series of blocks. However, it should be understood and appreciated that the claimed subject matter is not limited by the order of the blocks, as some blocks may occur in different orders and/or concurrently with other blocks from what is depicted and described herein. Moreover, not all illustrated blocks may be required to implement the methodologies described hereinafter.

[0031] FIG. 2 illustrates a flowchart of a method **200** for efficiently uploading data to a database using dynamic schemas. The method **200** will now be described with frequent reference to the components and data of environment **100**.

[0032] Method **200** includes an act of accessing one or more rows of feed data, the feed data including one or more data fields of fixed and varying lengths (act **210**). For example, data accessing module **110** may access feed data **101**. Feed data **101** includes multiple different rows of data, including row **1** (**104A**), row **2** (**104B**) and so on to row **N** (**104N**). As will be understood, substantially any number of rows may be included in feed data **110**. The data may include fixed-size fields **102** or variable-size fields **103**. Any given feed row data may include only fixed-size fields, only variable-size fields, or combinations of both. The feed data is represented by various extensible markup language (XML) files which include a definition of each field including the field's data type and data length. This field definition may be used by the data converting module **115** to convert the feed data to binary blob data **116** using a dynamic schema.

[0033] Method **200** includes an act of converting the received feed data into one or more portions of binary blob data, wherein the binary blob data is allocated in fixed-size data portions representing feed rows, and wherein varying length fields are represented by both a fixed-size field and a reference into the blob data that includes the actual varying size feed data (act **220**). For example, data converting module **115** may convert received feed data **111** into various portions of blob data **116**. The binary blob data may include two different portions. The first portion **117** includes fixed-size binary representations of each row, while second portion **118** contains data of varying-size fields. The first portion may contain references to the second portion for varying-size fields.

[0034] In some cases, the binary blob data representing multiple feed rows is sent to the database server **125** in a single round trip. Thus, rather than requiring one round trip for each row or table that is to be updated, all of the blob data may be sent to the database in a single round trip. The same technique may be used for converting and uploading data to a database server regardless of a particular feed data's schema and sizes. Still further, various portions of field data for fields with varying data size may be different in size for the same field within a feed data reading session. Thus, at any given

data reading, variable-size field data may differ in size for the same field, allowing for great flexibility in storing data in a database table data.

[0035] Method 200 includes an act of sending the binary blob data to a database server, the database server being configured to access metadata associated with a dynamic server statement to determine how to convert the binary blob data to a server table with a BLOB column configured to store the rows of feed data (act 230). For example, application server system 105 may send binary blob data 116 to database server 125. The database server accesses metadata 120 which includes a dynamic server statement. The metadata indicates to the server how the binary blob data is to be converted to a server table with a blob column 141 configured to store the rows of feed data 101. Execution of the dynamic server statement 121 (e.g. by module 135) then converts the data and stores it in data store 140.

[0036] In some cases, the application server may further determine that the binary blob data 116 is to be divided into data chunks. Each data chunk is uploaded to the database server 125 in one round trip. The application server determine that the binary blob data dividing is to occur based on one or more conditions. The conditions may include any one or more of the following: the number of rows in the binary blob data and the data size of the binary blob data. If one or more of the condition occurs (e.g. the blob data has a certain number of rows, or the data size is sufficiently large), then the data will be divided in a predetermined manner based on the condition that occurred. Still further, a database server presentation of the feed data is generated, where each feed row in feed data 101 is represented in a database table by one row with a single BLOB column containing feed row binary data (e.g. BLOB column 141 in data store 140).

[0037] FIG. 3 illustrates a flowchart of a method 300 for efficiently uploading data to a database using dynamic schemas. The method 300 will now be described with frequent reference to the components and data of environments 100 and 500 of FIGS. 1 and 5, respectively.

[0038] Method 300 includes an act of receiving from an application server one or more portions of binary blob data at a database server (act 310). For example, data receiving module 130 may receive binary blob data 116. This binary blob data may include fixed-size binary representation 117 and varying-size data fields 118.

[0039] Method 300 includes an act of executing a dynamic server statement to store the binary blob data into a database table, wherein the dynamic server statement is stored within metadata associated with the original feed (act 320). For example, the same technique may be used for converting and storing feed binary data in the BLOB column of a database table regardless of a particular feed data schema. In this manner, using of the static schemas may be avoided, as the schemas (or dynamic server statements) used are dynamic. Moreover, because the blob data is processed substantially entirely in memory, the data processing is performed in a highly efficient manner.

[0040] Method 300 further includes an act of restoring one or more feed data rows from the binary blob data by substituting varying size references with the actual feed data, while keeping the feed data in a binary format (act 330). For example, database server 125 may restore feed data rows 136 from the binary blob data by substituting different size references with the feed data 101, while keeping the data in a

binary format. Maintaining the data in a binary format allows the blob data to be processed by the database server in a highly efficient manner.

[0041] Method 300 also includes an act of storing the restored feed data in the database table in a blob data column, wherein each feed data row results in one binary blob value in one corresponding row in the database table (act 340). For example, database server 125 may store the restored feed data rows 136 in data store 140 in BLOB column 141.

[0042] FIG. 4 illustrates a flowchart of a method 400 for efficiently processing data using dynamic schemas. The method 400 will now be described with frequent reference to the components and data of environments 100 and 500 of FIGS. 1 and 5, respectively.

[0043] Method 400 includes an act of accessing feed data belonging to a particular feed (act 410). The database server then executes a dynamic server statement 510 to create a relational dataset with data type fields from the feed data in an in-memory table 515 of the server. The dynamic server statement is stored within metadata 120 associated with the feed data (act 420). The feed data is processed in the memory of the server, thus not requiring a corresponding static table, to aid in processing efficiency.

[0044] Method 400 further includes an act of applying a second dynamic server statement 520 to the feed data which applies various data processing conditions indicated in the metadata 120. The dynamic server statement is stored within the metadata associated with the feed (act 430). Then, upon determining that one or more feed data rows in database table 500 matches the data processing conditions, the feed data row information about the matching data rows is placed into an alert table that includes references to the original database table with feed blob data in BLOB column, thereby triggering an alert (act 440). Items placed in the alert table thus automatically issue an alert 521 to a user 525 that has subscribed to receive those alerts. Thus, the user can be apprised upon the occurrence of any of a number of different user-specifiable conditions. Moreover, referencing feed blob data from the original database table avoids duplicate data storing and using static representations of the feed fields.

[0045] Accordingly, methods, systems and computer program products are provided which efficiently upload data to a database using dynamic server statements. Moreover, methods, systems and computer program products are provided which efficiently process and store data using dynamic server statements. Such uploading, processing and storing may lead to a more efficient data management system.

[0046] The present invention may be embodied in other specific forms without departing from its spirit or essential characteristics. The described embodiments are to be considered in all respects only as illustrative and not restrictive. The scope of the invention is, therefore, indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

I claim:

1. A method, implemented at a computer system that includes one or more processors, for efficiently uploading data to a database using dynamic schemas, the method comprising:

an act of receiving from an application server one or more portions of binary blob data at a database server;
an act of executing a dynamic server statement to store the binary blob data into a database table, wherein the

- dynamic server statement is stored within metadata associated with the original feed;
- an act of restoring one or more feed data rows from the binary blob data by substituting varying size references with the actual feed data, while keeping the feed data in a binary format; and
- an act of storing the restored feed data in the database table in a blob data column, wherein each feed data row results in one binary blob value in one corresponding row in the database table.
2. The method of claim 1, further comprising an act of accessing the metadata included in the original feed to determine whether data processing conditions exist for the feed data.
3. The method of claim 2, further comprising, upon determining that one or more data processing conditions exist for the feed data, an act of creating a reference for each data row that satisfies the data conditions in a separate database server table.
4. The method of claim 3, wherein the reference includes a pointer to the feed data's original database server table that includes the associated binary feed data.
5. The method of claim 3, further comprising issuing an alert to one or more users based on a determination that at least one of the rows of the binary blob data satisfies the data conditions.
6. The method of claim 3, wherein the blob data is opened and processed in memory and stored in binary format, such that storing of static schemas is avoided.
7. A computer program product for implementing a method for efficiently uploading data to a database using dynamic schemas, the computer program product comprising one or more physical storage devices having stored thereon computer-executable instructions that, when executed by one or more processors of the computing system, cause the computing system to perform the method, the method comprising:
- an act of receiving from an application server one or more portions of binary blob data at a database server;
 - an act of executing a dynamic server statement to store the binary blob data into a database table, wherein the dynamic server statement is stored within metadata associated with the original feed;
 - an act of restoring one or more feed data rows from the binary blob data by substituting varying size references with the actual feed data, while keeping the feed data in a binary format; and
 - an act of storing the restored feed data in the database table in a blob data column, wherein each feed data row results in one binary blob value in one corresponding row in the database table.
8. The computer program product of claim 7, further comprising an act of accessing the metadata included in the original feed to determine whether data processing conditions exist for the feed data.
9. The computer program product of claim 8, further comprising, upon determining that one or more data processing conditions exist for the feed data, an act of creating a reference for each data row that satisfies the data conditions in a separate database server table.
10. The computer program product of claim 9, wherein the reference includes a pointer to the feed data's original database server table that includes the associated binary feed data.
11. The computer program product of claim 9, further comprising issuing an alert to one or more users based on a

determination that at least one of the rows of the binary blob data satisfies the data conditions.

12. The computer program product of claim 9, wherein the blob data is opened and processed in memory and stored in binary format, such that storing of static schemas is avoided.

13. A computer system, comprising:

one or more processors; and

one or more computer-readable media having stored thereon computer-executable instructions that, when executed by one or more processors of the computing system, cause the computing system to perform a method for efficiently uploading data to a database using dynamic schemas, the method comprising:

an act of receiving from an application server one or more portions of binary blob data at a database server;

an act of executing a dynamic server statement to store the binary blob data into a database table, wherein the dynamic server statement is stored within metadata associated with the original feed;

an act of restoring one or more feed data rows from the binary blob data by substituting varying size references with the actual feed data, while keeping the feed data in a binary format; and

an act of storing the restored feed data in the database table in a blob data column, wherein each feed data row results in one binary blob value in one corresponding row in the database table.

14. The computer system of claim 13, further comprising an act of accessing the metadata included in the original feed to determine whether data processing conditions exist for the feed data.

15. The computer system of claim 14, further comprising, upon determining that one or more data processing conditions exist for the feed data, an act of creating a reference for each data row that satisfies the data conditions in a separate database server table.

16. The computer system of claim 15, wherein the reference includes a pointer to the feed data's original database server table that includes the associated binary feed data.

17. The computer system of claim 15, further comprising issuing an alert to one or more users based on a determination that at least one of the rows of the binary blob data satisfies the data conditions.

18. The computer system of claim 15, wherein the blob data is opened and processed in memory and stored in binary format, such that storing of static schemas is avoided.

19. The computer system of claim 13, wherein the binary blob data represents feed data that includes a plurality of rows, each row including one or more data fields that correspond to varying length data, such that at least one row of the feed data has a varying data length relative to at least one other row of the feed data, and wherein the binary blob data comprises:

a first portion that represents each row of the feed data as a corresponding fixed-size row representation; and

a second portion which includes any varying length feed data.

20. The computer system of claim 19, wherein each corresponding fixed-size row representation includes (i) for any data field in the corresponding row of the feed data that corresponds to fixed length feed data, one or more corresponding fixed-size data fields; and (ii) for any data field in the corresponding row of the feed data corresponding to varying length feed data, both a fixed-size data field specifying a

length of the varying length feed data and a fixed-size data field specifying a reference into the second portion of the binary blob data, the reference identifying the actual varying length feed data.

* * * * *